ESTHER DYSON'S MONTHLY REPORT

Release 1.0®

VOLUME 19, NO. 10 | 20 NOVEMBER 2001 | www.edventure.com

Open Spectrum: The Paradise of the Commons

INSIDE

OPEN SPECTRUM 1
Fighting the Power 2
Box: Spectrum regulation
Computing Kills the Radio Star 5
The bandwidth illusion
Spread spectrum
Second coming of the Internet
802.11: The Seeds of Change 9
Learning to WiFi
Access for everyone
The grassroots
Room to Roam 14
NoCat.org
Commercial roaming services
The Open Spectrum Vision 18
Radio renegades
Spectrum politics
Commercial Meshed Wireless 21
Nokia Rooftop
UltraDevices
Spectrum Futures 25
Hardware to the rescue?

NEW MEDIA'S NEXT WAVE 27

Resources & Contact Information 30
Calendar of High-Tech Events 31

2002 PC Forum invitations are enclosed for subscribers! See back page for more on the event.

BY KEVIN WERBACH

Connectivity is a scarce resource. Building networks is difficult and expensive, something only a few companies can manage – telcos, cable operators, wireless carriers and satellite providers. Service providers control the fate of the Internet, as the gatekeepers to richer connections. The last mile into homes has become a broadband bottleneck, while high-speed wireless data services remain more wish than reality.

It doesn't have to be that way.

What if the scarcity and specificity of today's bandwidth options were artificial? What if wide-area network (WAN) connectivity were like that in local-area networks (LANs): a hardware choice under the control of the end-user? What if connectivity were shared by all rather than provided by a few?

It could happen. The new paradigm is open spectrum, and it requires only a few government decisions to open the floodgates for innovation. The rapid spread of 802.11b wireless LANs hints at the possibility of a wireless revolution as far-reaching and unexpected as the Internet was. So far, though, no government has considered the idea seriously. Open spectrum's disruptive potential is a political liability. If the technology industry wants to change the connectivity landscape for the better, now is the time to get involved.

Here's the concept in a nutshell: Instead of radio frequencies assigned exclusively to companies, spectrum would become a commons shared among users. Smart devices subject to rules ensuring that no one player could hog the airwaves would replace networks defined by governments and service providers. Bandwidth would be

{ *continued on page 2* }

cheaper and more ubiquitous. Spectrum would be used more efficiently. The gating factor for innovation would be vendors' ability to improve over time. . . and Moore's Law has long been a safe bet.

Fighting the Power

Broadband dreamin'

Talk of cheap, open broadband connections may sound unrealistic amid the current industry downturn. After pouring tens of billions of dollars into new entrants that collapsed under high debt loads and unrealistic demand projections, Wall Street is intensely skeptical of the communications infrastructure sector.

That's exactly the point. We're talking about wireless systems here, not the cable modem and digital subscriber line services that dominate today's high-speed access market. Open spectrum cuts out the greatest expense of communications networks: transport infrastructure. The network *is* the endpoints; in the middle is just air.

To make this approach possible, government agencies that control spectrum allocation should take four steps:

- Designate significant spectrum blocks for "unlicensed" use, meaning that anyone can transmit.
- Create equipment-certification rules to ensure transmitters in the new unlicensed bands cooperate intelligently for maximum efficiency.
- Allow low-power ultra-wideband services to "overlay" on licensed bands.
- Promulgate an "intelligent radio bill of rights" (an idea originated by Bran Ferren of Disney and now Applied Minds) to define the privileges and responsibilities of smart devices communicating over the airwaves.

The new rules would have to be clear and stable enough to give investors confidence to fund further development of unlicensed wireless technologies. Any of these changes would make innovative

Release 1.0® (ISSN 1047-935X) is published monthly except for a combined July/August issue by EDventure Holdings Inc., 104 Fifth Avenue, New York, NY 10011-6987; 1 (212) 924-8800; fax, 1 (212) 924-0240; www.edventure.com. It covers software, the Internet, e-commerce, convergence, online services, Web applications, technology policy, groupware, data networking, streaming media, enterprise applications, wireless communications, intellectual property and other unpredictable topics.

EDITOR-IN-CHIEF: Esther Dyson (edyson@edventure.com)

PUBLISHER: Daphne Kis (daphne@edventure.com)

EDITOR: Kevin Werbach (kevin@edventure.com)

SYSTEMS DIRECTOR: Scott Giering (scott@edventure.com)

CIRCULATION MANAGER: Natasha Felshman (natasha@edventure.com)

OFFICE MANAGER: Bill Demeritt (bill@edventure.com)

EXECUTIVE ASSISTANT: Irene Lawrence (irene@edventure.com)

CONSULTING EDITOR: Bill Kutik (bill@kutik.com)

Copyright © 2001, EDventure Holdings Inc. All rights reserved. No material in this publication may be reproduced without written permission; however, we gladly arrange for reprints, bulk orders or site licenses. Subscriptions cost $795 per year in the US, Canada

new services possible; all of them together would lay the groundwork for a new era of communications.

Open spectrum wouldn't break the bandwidth bottleneck overnight. In practice, there will still be limits on how many users can communicate effectively, depending on available frequencies, power, competing uses and the design of transmitters and receivers. The argument for open spectrum is that a spectrum commons model is *more* efficient than the traditional licensing model, and that gap will widen over time. Moreover, the airwaves are a public resource. Giving everyone the ability to transmit would promote democratic values and stimulate user-generated innovation, just as the Internet did.

A true spectrum commons can't emerge today because government rules preclude it. Under those rules, only spectrum licensees can transmit in most frequency bands. And with no market, few investors or researchers are working on the underlying technologies. Regulatory uncertainty has hamstrung startups such as Time Domain and Fantasma Networks working on ultra-wideband (UWB), one form of unlicensed wireless technology. Fantasma, a spinoff of Interval Research, shut down when the US government moved too slowly to authorize UWB services.

As a matter of fundamental physics and information theory, we simply don't know the limits on how efficient a spectrum commons could be. But we can't afford not to find out. As Paul Baran, the father of packet switching, has observed, "Any spectrum space not being used is water over the dam – and forever wasted."

Breaking the vicious circle

Breaking out of this vicious circle won't be easy. The companies who paid dearly to control scarce spectrum oppose alternatives that would make spectrum shared and virtually free. They argue that expanding unlicensed services would cause ruinous interference, or what ecologist Garrett Hardin termed a tragedy of the commons. The real tragedy is otherwise. The individual interests of today's spectrum owners are preventing a commons that could benefit all.

Two developments are pushing toward open spectrum. In the slivers of spectrum where shared use is allowed, a new industry is emerging around the 802.11 wireless LAN standard (SEE PAGE 9). 802.11 is moving from a technology to connect home or office PCs to a widespread mechanism for high-speed public Internet access. There will be 4,000 wireless access points in locations such as airports and cafes by the end

DIVVYING UP THE AIR: A BRIEF INTRODUCTION TO SPECTRUM REGULATION

Spectrum regulation began early in the 20th century with the commercial development of radio. In the US, after a period of rapid technical development, bitter patent battles and fights about interference, the federal government passed the Radio Act of 1927. This legislation, the precursor to the Communications Act of 1934 which established the Federal Communications Commission (FCC), gave the government the authority to assign radio spectrum bands to broadcasters, and to prevent others from interfering with transmissions in those bands.

Spectrum delegation is a two-step process. First government *allocates* spectrum to particular uses and defines requirements such as power levels and fixed or mobile usage. Allocation involves many tradeoffs. For example, should a mobile phone service be designed with one or two national licenses, or dozens of local ones? After making those decisions, government *assigns* that spectrum to private entities, typically through perpetual licenses. In other words, the company doesn't own the spectrum, which is considered a public resource.

Government retains large swaths of spectrum for military and other public uses, and defines some spectrum as "unlicensed" and available for applications ranging from ham radio to cordless telephones. To prevent interference, government sets rules for transmitters and receivers, though today much of that process occurs through self-certification.

In the US, the FCC is responsible for commercial spectrum bands and the National Telecommunications and Information Administration within the Department of Commerce oversees government spectrum. The little-known Interdepartment Radio Advisory Committee, with membership from all relevant federal agencies, reviews proposed spectrum allocation changes. Other countries have similar regimes, though the exact spectrum assigned to particular uses varies.

The spectrum itself doesn't change, but uses of spectrum evolve over time. As technology improves, it makes some applications possible in bands where they previously were infeasible. Some applications such as mobile telephony weren't commercially viable in earlier decades. As uses and priorities change, government alters the rules, and sometimes relocates existing users to clear out space for other uses.

The biggest change in the past decade has been the method of spectrum allocation. Since there isn't enough spectrum for everyone, how to choose among competitors? Originally the method was competitive hearings or outright grants to dominant companies such as incumbent local telephone carriers. This was inefficient and was seen by many as unfair, so the FCC switched to lotteries for early mobile phone licenses. The result was further chaos. Individuals signed up and won licenses which they immediately resold to companies actually interested in building the systems.

Finally, in the early 1990s, the FCC began auctioning spectrum. The rationale behind auctions is that the companies who value the spectrum most will pay the most for it, and having done so, they will have strong incentives to build services rather than let the spectrum lie fallow.

Beginning with second-generation mobile phone licenses, spectrum auctions have generated hundreds of billions of dollars in revenue for the US and other governments. However, the high costs of auctioned spectrum have come under fire. Some auctions have failed when winners went bankrupt, leading to litigation.

of this year, and 38,000 in 2006, estimates Gartner Dataquest. The second open-spectrum thrust comes from a loose group of academics and technologists (SEE PAGE 18). They are promoting the idea of a spectrum commons to policy-makers and researching the engineering foundations for shared spectrum use.

Unlicensed wireless services will first take hold at the grassroots level, among non-traditional service providers and companies targeting markets underserved by DSL and cable modem service. In this way, 802.11 may become a stealth competitor for last-mile residential broadband. "It's very much like the early days of the ISPs. What's invisible to most VCs and analysts is this huge groundswell that's just below

the radar. It will look like it suddenly popped out," says Robert Berger, founder and chairman of meshed routing vendor UltraDevices (SEE PAGE 23).

The hope is that, as these pockets link up and become noticed, pressure will coalesce for broader application of shared-spectrum approaches. The danger is that 802.11 is a poor standard-bearer for open spectrum. The fact that some unlicensed services are possible may alleviate pressure to open up room for more.

Computing Kills the Radio Star

The bandwidth illusion

The concepts behind open spectrum were expressed long ago. Spread spectrum, its technical foundation, is based on work done in the 1940s.[1] The view that spectrum could be managed more efficiently by private actors than by governments dates from the 1950s. Proposals for shared spectrum were floated in the 1980s by the US Federal Communications Commission (FCC) itself. [DISCLOSURE: KEVIN WERBACH WORKED AT THE FCC BETWEEN 1994 AND EARLY 1998.] However, those proposals were abandoned. No industrialized nation has taken open spectrum seriously because its ideas are so dangerous and counter-intuitive.

New York University law professor Yochai Benkler (SEE PAGE 18) summarizes open spectrum's most startling claim: "Bandwidth is just a parameter in an equation; it's not a naturally bounded resource." We think of the spectrum as a physical entity that can be subdivided and arranged on a dial. In reality, there's no there there. The only elements of any radio communications system are a transmitter, a receiver, and an electromagnetic signal between them. And the signal doesn't really travel from one specific point to another; it radiates outward indefinitely, until it attenuates below the threshold of detection.

The magic happens at the endpoint devices. When your mobile phone or wireless email device receives a message, what it's doing is picking out one signal from all the radio waves washing over it, and disentangling that signal from background noise. My noise is your signal. The radio waves themselves pass through each other; what

[1] Direct-sequenced spread-spectrum services distribute coded signals across a range of frequencies to receivers that scan for them and reassemble the message. Frequency-hopping spread-spectrum radios shift signals between different bands, further avoiding congestion. Ultra-wideband services use such low power over such a wide swath of spectrum that to traditional receivers their signals are indistinguishable from background noise, potentially allowing them to co-exist with licensed services in the same bands.

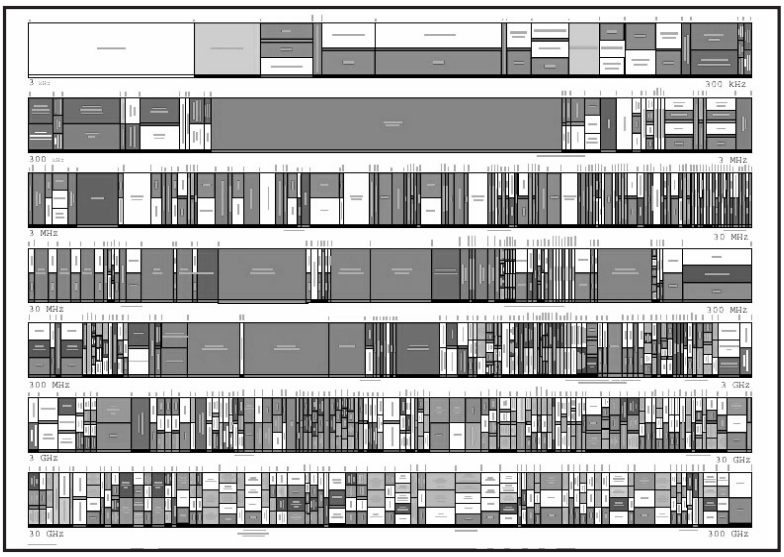

FIGURE 1: *US spectrum allocation table.*

we call "interference" is a function of each receiver's ability to distinguish among them. Wireless communication is inherently an exercise in computation.

In fact, this is true of any communication. Bell Labs researcher Claude Shannon, who developed information theory in the late 1940s, showed mathematically how systems could encode and decode signals to send information efficiently. (Along the way he invented the concept of a bit.) Though wired and wireless communications are regulated in totally different ways, the only technical difference is the physical media they transmit over. With no wires to contain signals, wireless devices must work harder to separate distinct communications.

If Al wants to send a wireless message to Ethel, he can transmit it with high power, so that it's easy for Ethel to distinguish it from background noise. But the higher the power, the more other signals are drowned out. Everyone else has an incentive to use high power to be heard, resulting in congestion. Imagine trying to understand someone across a crowded room in which everyone is shouting at the top of their lungs.

Historically, the solution to this problem was to split up the spectrum into narrow bands and assign them exclusively to providers in a defined region (SEE FIGURE 1). With power and usage restrictions, this meant that many users could share the airwaves. A monopoly within each band led to some degree of diversity overall. To continue the previous analogy, only one person would be allowed to talk at a time, or each person would be allowed to talk in a particular part of the room.

Start spreading the bits. . .

The conventional solution to the wireless tragedy of the commons depends on assumptions about technology. If the receivers are smart enough, many broadcasters can send signals with *low* power over a wider range of bands and not interfere. So everyone in the room can talk at once, so long as they do so quietly. And if you listen carefully, you'll recognize the unique pitch of a friend's voice across the room despite the distance and the large number of simultaneous conversations.

This approach is called spread spectrum. In a spread-spectrum system, low-power signals are distributed across many frequency bands and then reassembled on the other end. Content-delivery networks such as Akamai do something similar to speed up Websites: trade computation for bandwidth. It's faster and cheaper today to use algorithms to locate the closest copy of content on a local caching server than to deliver that content over a big pipe from the origin.

Spread spectrum is widely employed in mobile telephony, radar, satellite and military networks. For example, Qualcomm's code division multiple access (CDMA) technology, the basis for many of the world's cellular networks, uses spread spectrum techniques. But this approach is more than a mechanism to use licensed spectrum efficiently; it challenges the need for licensing to begin with. When spread spectrum was developed, there wasn't sufficient computing power to take it to its logical conclusion. Now, thanks to advances in digital signal processing and other technologies, the balance has shifted.

Spread spectrum by itself isn't enough. To realize the open spectrum vision, transmitters and receivers must cooperate using smart "software-defined" radios. Nodes would sometimes act as repeaters for traffic between other nodes; they would dynamically select power levels and coding schemes based on observation and cooperation with other nodes; and they would cooperatively sample background noise and the propagation environment, and would adapt based on that sampling.

It's the equivalent of what multiprotocol routers do today to manage traffic on the distributed Internet. Explains David Reed, one of the original developers of Internet protocols and formerly chief scientist of Lotus: "Each of these techniques has been shown to dramatically change the capacity of the network as a function of the number of nodes and amount of traffic. They amount to distributed algorithms that provide an overall speedup that grows as the number of nodes increases." In other words, more transceivers make congestion less of a problem, not more.

The most detailed technical analysis of how large numbers of users could share unlicensed spectrum for high-bandwidth data communications is Tim Shepard's 1995 MIT doctoral dissertation (SEE RESOURCES). Shepard and Reed are among the few engineers exploring this approach today. Like the packet-switched data networking research in the 1970s that produced the Internet, open spectrum is outside the mainstream today, but could deliver enormous benefits tomorrow.

The second coming of the Internet

A spectrum commons wouldn't just work like the Internet; it would catalyze innovation for the same reason the Net has. Traditional communications networks are based on expensive pipes and cheap endpoints. The pipes are the bottleneck: if everyone had unfettered access, they would quickly overload.

The Internet flipped the model, using intelligence at the endpoints of a radically simplified network (SEE **RELEASE 1.0**, DECEMBER 1999). Innovation flourished when it depended on routers from Cisco and software from Microsoft, AOL and Yahoo! instead of network-wide infrastructure upgrades by phone companies. Government played a key enabling role by preventing telcos from limiting what could be connected or pricing Internet service providers out of the market.

Open spectrum means doing the same thing in wireless. It means moving from centralized networks built by carriers to decentralized services based on smart transmitters and receivers. There's one big difference: Spectrum is everywhere. A true wireless Internet is a universal Internet, not dependent on any infrastructure provider. And it's a broadband Internet to boot.

The dominant wireless model is hitting a wall. Scarce resources inevitably run out, and as they approach that point they become more expensive. Companies are fighting over ever-smaller scraps of spectrum; hence the budget-busting billions European carriers committed for third-generation (3G) mobile phone licenses, and the legal battle in the US over wireless spectrum won by bankrupt NextWave. In this environment, governments and industry need to consider alternatives.

Signal to noise

Opponents of open spectrum raise various arguments to the effect that there's no such thing as a free lunch. Systems using smaller cell sizes and shorter links can pack more users into the same area, but the additional hops can create unacceptable latency for services that must travel long distances. However, that isn't a reason to prevent experimentation. Latency and other effects depend on the architecture of the system and how extensively radios cooperate.

Dale Hatfield, who until earlier this year served as the FCC's chief technologist, agrees on the need to consider decentralized approaches. "I think it's important that we explore all these avenues and try this stuff," he says. But he doesn't believe open-spectrum advocates have proven their case yet: "It's not clear that when you're all

done, you can fit everyone in with an adequate quality of service, while still allowing other systems that probably still have to be done on a centralized basis."

Those other services include Global Positioning System (GPS) satellite-based location, air-to-ground communications with airplanes and radios used by public safety organizations such as police and fire departments. Protecting or relocating those applications puts you squarely back in the realm of FCC allocation procedures.

Users of these services raise legitimate concerns, especially as they touch on public safety. However, they are often the cause of the problem. The "interference" happens because receivers for services such as GPS aren't as intelligent as they could be in distinguishing signal from noise. This is a recurring issue. The standards for broadcast television, for example, required channels to be widely separated, effectively codifying 1950s technology. A standard written today would allow many more channels.

It's a classic economic externality. Just as a factory has no reason to stop dumping pollutants in a river unless it's made responsible for the costs it imposes on people living downstream, spectrum users need incentives to use receivers that increase spectral efficiency rather than reducing it.

Given the potential gains, we agree with Benkler that the onus should be on existing users, rather than on proponents of unlicensed spectrum, to justify themselves: "The debate should shift to how much spectrum we should just free, and what slivers – if any – we need to maintain as owned."

802.11: the Seeds of Change

Learning to WiFi

Fortunately, there is a test case for these views: 802.11 LANs. These aren't true distributed, cooperative networks, but they do use spread spectrum to provide high-speed data communications over unlicensed frequency bands. Thus, it's instructive that 802.11 has become the hottest technology in the networking world. As WiFinder chairman Scott Rafer (SEE PAGE 16) points out, "license-exempt market share capture is *de facto* deregulation of the wireless data business."

The first generation of 802.11 services use the 2.4-GHz unlicensed band, known officially in the US as ISM (institutional, scientific and medical) because of its origi-

nal uses. It's designated as unlicensed virtually everywhere in the world. There is also unlicensed spectrum in the US in the 900-MHz range (used by many cordless phones and by now-bankrupt Metricom's original low-speed wireless Internet access service), as well as the unlicensed national information infrastructure (U-NII) band in the 5-GHz range.

Such unlicensed bands were designed for short-range transmissions, which don't need much power. A cordless phone, for example, need only communicate with a nearby base station. Even without spread spectrum, someone across the street can use a remote control on the same frequency and not interfere with your reception. Having a carrier to sell end-users the spectrum between their cordless phones and base stations makes no sense. Conventional wisdom was that it would be hard to do anything particularly interesting at 2.4-GHz, because of the motley assortment of short-range services operating there. Hence its unofficial name: the "junk" band.

Despite this, several protocols have been developed in the past several years to use the 2.4-GHz band for wireless local-area networking. These include HomeRF, Bluetooth and 802.11b (SEE **RELEASE 1.0**, APRIL 1999). Cybiko (SEE **RELEASE 1.0**, MARCH 2000) uses a proprietary 2.4-GHz technology to connect users of its handheld gaming and communication devices. [DISCLOSURE: ESTHER DYSON IS ON THE BOARD OF CYBIKO.]

802.11b is the one generating the most excitement. 802.11 refers to a family of protocols standardized by the IEEE for wireless Ethernet communications. (Traditional wired Ethernet is 802.3.) These are also called WiFi systems, a marketing term coined by the Wireless Ethernet Compatibility Alliance (WECA), a trade association. Confusingly, the first variant widely deployed is 802.11b, which offers transmission speeds of 11 Mbps.

802.11b systems require only sub-$100 networking cards, a hub called an access point (now about $150) and an Internet connection. Starting with Apple's 1999 introduction of Airport wireless networking in its laptops, products based on 802.11b have sold briskly, with an estimated 10 million cards in use by the end of 2001. Says Berger of UltraDevices, "It's the Internet pushing out a new capillary."

A Microsoft-sponsored survey this summer of 180 companies with at least 500 PCs found that 40 percent were using wireless networks, and another 31 percent intended to within 18 months. Analyst firms predict the 802.11 equipment market will generate several billion dollars annually in both the US and Europe within five years.

One benefit of 802.11 systems is their flexibility. They are being deployed for home networking (linking PCs to share printers or Internet connections), small-office networks (as an alternative to running Ethernet cable), large businesses (offering workers freedom of movement around a building or corporate campus), universities (allowing students to be connected throughout the campus area) and public access points (SEE BELOW). Centerbeam, which bundles computers, software and Internet access for businesses for a monthly fee, uses 802.11b because it eliminates wiring as a point of failure. If a piece of equipment fails, Centerbeam just sends a new one.

802.11a units, which can send data at 54 Mbps, are now coming onto the market.[2] They transmit in the uncongested 5-GHz U-NII band. Dual-radio devices can support both protocols simultaneously, allowing for backward compatibility or 802.11a backbones connecting local 802.11b access points. In volume, 802.11a chipsets from companies such as Atheros have only a modest price premium over 802.11b.

Access for everyone

802.11 is fast becoming a core part of the computing ecosystem. Current-generation laptops often have 802.11 cards built in standard or as an option. Windows XP includes native support for the protocol, as does Apple's OS X. Many residential gateway devices (SEE **RELEASE 1.0**, JULY/AUGUST 2001) double as 802.11 hubs. Public wireless access "hotspots" are appearing everywhere, as are private networks. NYCWireless co-founder Terry Schmidt (SEE PAGE 13) recently found more than 4,000 wireless LANs by driving around Manhattan with a laptop running NetStumbler, a free program that identifies nearby wireless LANs.

Security remains a serious issue for the nascent wireless LAN industry. A large number of the access points Schmidt identified allowed open connections from anyone in range. (The practice of driving around looking for such open nodes is called "war driving," after the old "war dialing" method of calling phone numbers in sequence to find insecure dial-in connections to corporate networks.)

Even when access points are configured properly to limit access, they may not be secure. The standard security technology for 802.11b, Wireless Equivalent Privacy (WEP), is supposed to make wireless LANs as secure as their wired counterparts. However, WEP has been shown to be insecure and easily cracked. Newer access

2 The IEEE has defined several higher-speed 802.11 variants. There is also a competing technology, HiperLAN/2, which the European Telecommunications Standards Institute (ETSI) has endorsed, setting

points have stronger security built in, and the IEEE is now standardizing a better security mechanism called 802.1x. 802.1x also offers back-end authentication that can be used to tie together access points on different networks.

Because 802.11 is an industry standard, with good compatibility among vendors, pretty much any device can physically connect to any access point. A user who installs an 802.11 card for an office network can use it to tap into a wireless network at a hotel or a conference (including our PC Forum). In other words, 802.11 lends itself to service business models in addition to hardware sales.

Two venture-funded service providers, Wayport and MobileStar, are hooking up access points in high-traffic areas, giving users wireless connections in all their locations for a monthly fee or on a pay-as-you go basis. MobileStar, which is putting 802.11 connections into Starbucks cafes in partnership with Microsoft and Compaq, recently ran into financial trouble and laid off most of its employees. However, its network remains operational and it has reached a tentative agreement to sell itself to VoiceStream Wireless.Wayport has focused on hotel chains and airports.

MobileStar and Wayport have received the most attention and funding, but they aren't the only commercial 802.11b services providers. For example, Surf and Sip, a four-person boot-strapped startup, has established access points in cafes and coffee houses in San Francisco; Los Angeles; Colorado Springs; Fish Creek, WI; Fort Lauderdale; and Jackson Hole. Companies that provide high-speed wireless connections to hotels such as STSN are also getting into the 802.11 game.

The grassroots

With 802.11 networks, there can be an Internet access service without a service provider. The great limitation for high-speed Internet connectivity today is the last mile (SEE **RELEASE 1.0**, JULY/AUGUST 2001). There is plenty of backbone fiber in populated areas, but few affordable options for getting it to end-users. 802.11 is perfectly suited for this scenario. It's a short-range technology with very low costs, because it requires only edge hardware.

Technically savvy 802.11 networking users have made this connection and have established public wireless access points. This is occurring at the same time as commercial fixed-wireless Internet access providers such as AT&T and Sprint pull back. "Insofar as wireless networks take off at the consumer level it will be because of grassroots efforts," predicts Schuyler Erle of NoCat.org (SEE PAGE 15).

This isn't so far-fetched. You may notice that the cable TV jack on your television set says CATV; that stands for community antenna TV. The original cable systems were established by communities with poor over-the-air TV reception. They set up large central antennas and shared them by running coaxial cable to households. To oversimplify a long story, the antenna providers built networks, consolidated and started using satellite and fiber-optic connections to deliver programming to their systems, becoming the commercial cable operators we know today.

Community wireless networking groups have sprung up in numerous metropolitan areas around the world. These groups set up public 802.11b access points in neighborhoods, parks and coffee shops which members or the general community can use for free. The list maintained by the Personal Telco Project in Portland, OR (SEE RESOURCES) includes more than 80 operational community networks. The wireless community networking scene today calls to mind the early days of ISPs, especially the Freenet movement that provided connectivity to communities before commercial Internet access was available.

New York's NYCWireless is a typical example of an 802.11 community networking group. The organization was formed in May by Terry Schmidt, an independent consultant, and Anthony Townsend, an urban-planning graduate student. Schmidt established the first NYCWireless access point in an Upper East Side coffee shop. Townsend set up a node in Washington Square Park using his network connection at New York University (NYU), where he is a fellow at the Taub Urban Research Center. NYCWireless now sponsors fifteen access points around New York City.

NYCWireless got a boost following September 11. Thousands of businesses near the World Trade Center were displaced, and wireline networks were disrupted when the Twin Towers collapsed. NYCWireless set up temporary point-to-point wireless links for Pace University, connecting their internal networks to functioning Internet access points elsewhere. To help with these and other community-oriented efforts, the WECA trade association donated $30,000 worth of equipment to the group.

The group hopes to address potential legal issues by requiring users to agree to an acceptable use policy when they first connect. A participant is writing a Java application that identifies all current users on NYCWireless access points throughout the city, potentially forming the basis for new kinds of applications that meld the physical and the virtual worlds (SEE **RELEASE 1.0**, MAY 2001 AND JUNE 2000).

"People are interested in being the networking equivalent of the open-source movement," says Townsend. Like open source, though, community networks have been attacked by those who see them cannibalizing commercial businesses. The local 802.11 traffic uses free unlicensed spectrum, but to get the bits onto the Internet requires a traditional outside connection.

Public access points are often connected to residential DSL or cable modem accounts, or to business networks. Managers of such networks frown on users making network connections freely available to the general public, which may cause congestion and open them up to security risks. Townsend, for example, had to shut down the original Washington Square NYCWireless node at the behest of NYU, which didn't want its campus network used for public access.

Seeing parasitic use of commercial networks as the primary function of community wireless access points misses the point. Node operators are subject to the usage policies of the networks they subscribe to. Mechanisms are being developed to ensure they don't overload those networks or create security holes. There are many ways to fund community access points, from co-operative group membership to public subsidies by local organizations to locations such as cafes paying for the outside bandwidth to attract more customers.

Room to Roam

802.11 wireless networks are exciting, but their benefits are limited if users can connect to only one access point. Because 802.11 is a LAN technology, it was designed with local connectivity in mind. Users must authenticate themselves each time they connect; using a new network can require setting a parameter called the SSID and having a local account. Windows XP and Apple's OS X make connecting to 802.11 networks much simpler than before, but authentication can still be a problem.

Thus, roaming is essential for 802.11-based services to go beyond niche offerings in local hotspots. "In a license-exempt environment, no single network operator will achieve ubiquity for a decade," notes WiFinder's Scott Rafer.

In the dial-up Internet world, large ISPs were able to build out sufficient local footprints to make roaming necessary only for frequent travelers, especially overseas. That's not going to happen in wireless, because there are too many different kinds of

locations to cover for a national 802.11 ISP to make sense, especially in this difficult capital-raising environment.

Dial-up Internet access and mobile phones both use area codes as the smallest local unit, since calls within an area code are generally unmetered in the US. Most usage happens with a local call. The reverse is true for 802.11 services. "Most access instances are roaming, because of the nature of the network. It's very fragmented and it's very small," explains HereUare ceo Clark Dong (SEE PAGE 16).

Wireless roaming will be achieved in several ways. We describe below some companies and community groups attacking the problem. WECA, the 802.11 trade association, is also working with vendors to promote interoperable authentication.

Letting the roaming cat out of the bag

NoCat is a group in Sebastopol, CA, made up primarily of employees at technology publisher and conference organizer O'Reilly & Associates. The name comes from a description of radio by Albert Einstein. It's like a very long cat, he observed: you pull its tail in one place and the head meows in another place. . . only there's no cat! NoCat operates community wireless access points in Sonoma County, CA. It is also developing open-source software and an authentication service to allow roaming across public wireless nodes.

As Schuyler Erle, lead developer for the project, explains: "The authentication software was devised in an attempt to address the all-or-nothing security issues of 802.11b as it originally shipped. You have a choice of either having an open access point – in which case you're giving all the bandwidth you have to everyone – or you run a closed access point – in which case you can't share your bandwidth with your neighbors and your community."

NoCat is writing authentication software that allows a member of a community wireless group to authenticate at an access point run by another group. It supports differentiated quality of service, so that an access point operator can provide best-efforts service to anyone and higher-quality connectivity to members. This approach also allows the operator to limit overall bandwidth usage, so that the access point doesn't overload the rest of its network. "It was our idea that each local cooperative would set its own standards," says Erle. NoCat has set up a central authentication server that other groups can use if they don't want to take on the cost and management hassles of running one.

Commercial roaming services

A number of commercial services are looking to facilitate wireless roaming. GRIC and iPass, the two main roaming networks for dial-up, now support 802.11 services. These companies have developed authentication, billing and mediation infrastructure that lets users of one ISP connect through the access network of another. For dial-up, the primary users are business travelers, especially those who travel internationally. As ISPs offer 802.11-based services, their subscribers can use the same GRIC and iPass back-end systems to roam across providers.

HereUare, a two-year-old, 50-person startup based in San Jose, CA, has raised $16 million in two funding rounds from ComVentures, Raza Foundries, Lucent Ventures and Pacific Ventures. "We're providing infrastructure services to enable anybody to become an operator, taking advantage of the fact that this wireless LAN technology is very cheap today," says ceo Clark Dong. He continues: "Because it's so cheap, you need a fabric underneath it all to stitch it together into one large network."

In addition to the authentication and mediation functions that GRIC and iPass offer, HereUare provides back-office services such as subscriber management, billing and provisioning to service providers. This allows anyone to deliver 802.11-based access, with roaming, just by purchasing an access point and a backhaul connection to the Internet.

HereUare has announced a relationship with MobileStar, and though he won't disclose specifics, Dong says the company is working with "all the major guys" in the 802.11 access space. HereUare is currently offering free roaming to generate interest in its services. It plans to charge service providers fees depending on the level of back-office services it provides. "In the future, consumers won't know what network they are on," says Dong.

In Europe, Luxembourg-based Excilan hopes to enable 802.11 roaming modeled on the European mobile-phone industry. It is creating a central clearinghouse to broker and settle between wireless service providers.

HereUare's closest competition in the US is WiFinder, which comes at the roaming question from a different direction. The startup, based in Lafayette, CA, builds coverage maps for public and corporate 802.11 access points that show the location of hotspots. After all, if you want to roam across wireless access points, you have to know where they are!

At WWW.WIFINDER.COM, users can type in an address and find nearby access points offered through Wayport, MobileStar, community access groups and other providers. As 802.11 access points spread, WiFinder's maps may help customers make decisions about which service provider to use, and may help those providers determine where new access points would be most valuable.

Unlike HereUare, WiFinder isn't involved in the back-end authentication and mediation processes for 802.11 node owners. Instead, its primary customers will be corporate IT departments. "Vendor bundling of 802.11b radios in many laptops in 2002 is very disruptive to security, network management, and remote access policies," says chairman Scott Rafer. Many large companies are deploying their own 802.11 nodes for employees. Others simply need to support mobile workers who take advantage of wireless hotspots when they travel. WiFinder will offer software and services to allow IT departments to integrate these new networks into their existing systems.

Send in the carriers

Unlicensed local-area wireless connections don't necessarily conflict with traditional mobile phone services. Despite the promise that 3G wireless services will deliver ubiquitous high-bandwidth wireless connections, the reality today is that wireless data services are unreliable and bandwidth-challenged. Even as operators spend billions deploying next-generation wireless data networks, they won't have universal coverage any time soon.

Why not combine 802.11 LANs and wide-area licensed wireless services, giving users the best of both worlds? That's exactly what's starting to happen in some areas. Nokia offers a roaming solution that combines an 802.11b transceiver card with a GSM mobile phone SIM card. When a user is within range of an 802.11 access point, he or she gets 11-Mbps local connectivity. Everywhere else, the system switches to the wide-area wireless data network.

Nokia claims that 20 mobile operators are testing or running pilots, and expects commercial deployment in the first half of 2002. Finland's Sonera announced its wGate service based on Nokia's technology earlier this month.

The Open Spectrum Vision

Radio renegades

While all this transpires in the marketplace, a small cadre of scholars is building the intellectual foundations for the open-spectrum concept. Like the Web, open spectrum has tremendous commercial potential that established businesses may not recognize. The notion of using smart devices instead of intelligent networks runs counter to communications industry received wisdom. Finally, the idea of a spectrum "commons" open to all takes some getting used to.

In the face of such skepticism, there are two overwhelming arguments for exploring open spectrum. The potential economic benefits of using spectrum more efficiently and facilitating innovation are huge. And the social benefits of wireless spectrum as an open communication medium may be greater still.

The leading open spectrum advocates are divided among policy advocates and technologists. Law professor Yochai Benkler penned a 1998 article in the *Harvard Journal of Law and Technology* questioning the rationale for licensing spectrum in the current technological environment. (We took the phrase "open spectrum" from the title of a workshop Benkler organized this spring.) Larry Lessig, well-known for his involvement with the Microsoft antitrust case and his analysis of the legal implications of software code, has also joined the cause. Dave Hughes of Old Colorado City Communications has long advocated unlicensed wireless systems as the best way to bring connectivity to rural communities.

On the technical side, Dewayne Hendricks, head of a wireless systems integration company and a former executive at cable modem vendor Com21, leads the spectrum working group of the FCC's Technological Advisory Council. David Reed, who, in addition to his other work, co-authored the "end to end" article that defined the Net's fundamental architecture (SEE **RELEASE 1.0**, MAY 1999 AND FEBRUARY 1999), is exploring the foundations for cooperative wireless networking.

"We could have the greatest wave of innovation since the Internet (and probably bigger in impact, because more pervasive) if we could unlock the spectrum to explore the new possibilities," says Reed, summarizing the case for open spectrum.

AN IMPLEMENTATION IS WORTH A THOUSAND ARGUMENTS

The best way to show the benefits of open spectrum, and to allay concerns, is to put it into effect. That can't happen on a wide scale in the US, though, because cooperative spread-spectrum is illegal in existing licensed bands and unlicensed bands are too crowded and restricted. So open-spectrum proponents are taking their efforts to other nations. . . some of them inside the US!

Dewayne Hendricks' Dandin Group is deploying unlicensed wireless communications systems in the South Pacific island nation of Tonga. Tonga is an interesting case; it also signed a deal a few years ago to make its .to top-level Internet domain name available to commercial interests as an alternative to .com. Hendricks is also talking with Native American tribes about deploying spread-spectrum communications services on reservations in the US. Tribes are sovereign nations under US law, and thus arguably not subject to FCC spectrum regulations.

Hendricks is keeping the US government abreast of his efforts, but he expects opposition from spectrum license holders as the project gets off the ground. Lessig has agreed to handle the legal defense work, should that become necessary.

Shifting the debate

What distinguishes the open-spectrum advocates isn't their view that current spectrum policy is broken. Even FCC Chairman Michael Powell agrees with that. "Our nation's approach to spectrum allocation is seriously fractured," he said at a press conference on October 23rd.

Everyone has his or her own reasons to hate the status quo. Wireless operators claim they can't get their hands on enough spectrum to meet demand. Free-market advocates criticize the FCC for its rules limiting how spectrum can be used and how much companies can own. Financial markets worry about the high prices carriers have paid to acquire spectrum in recent auctions. Different classes of users – mobile phone operators, television broadcasters, public safety agencies such as police and fire departments, military agencies, the aviation industry, satellite communications firms, ham radio operators and others – argue about who should get priority.

The open-spectrum proponents seek to change the terms of the argument. That's what makes it so hard for any of the existing players to grasp. Says Benkler: "The thing that is stickiest is the basic mindframe that comes from 50 years of structured conversation [with the choice] being between markets and licensing." Benkler's work explains how open spectrum policies would be at least as economically efficient as the current auction regime, and likely more so.

Spectrum politics

The FCC and regulators in other developed countries have so far been unwilling to adopt the open-spectrum approach. Under chairman Bill Kennard and chief technologist Dale Hatfield, the FCC did launch proceedings on ultra-wideband technology and software-defined radios, and officials met with open-spectrum proponents

on many occasions. With the involvement of Hendricks and others, the Technological Advisory Council analyzed issues related to spread spectrum.

In his October 23 press conference, current FCC chairman Michael Powell pledged to "continue strong support for unlicensed bands where feasible." His reference to unlicensed bands as a major element of spectrum policy was significant. On the other hand, he offered no proposals for expanding, as opposed to maintaining, the conditions under which shared spectrum is possible. The FCC played an underappreciated role in protecting early ISPs from regulation and threats by telcos; it could open up a similar space for non-traditional services in the wireless arena.

If rules governing technologies such as ultra-wideband continue to constrain the development of the market, the resulting services will not use spectrum as efficiently as they could if the engineers had free rein. For example, if 802.11 becomes "the" technology for unlicensed wireless data services because it is able to operate under the existing spectrum licensing regime, investment won't be directed towards technologies and architectures that might provide a more effective solution to the last-mile bottleneck.

In the 1980s and early 1990s, the champions of spread spectrum were conservative economists and futurists such as George Gilder. They saw these new technologies as a way around top-down command-and-control regulation by governments. The Reagan-era FCC launched a spread-spectrum proceeding in the mid-1980s, initially proposing to allow overlay of unlicensed spread-spectrum uses with licensed services. Gilder promoted Qualcomm's CDMA technology, which eventually became one of the major standards for second- and third-generation wireless phone systems.

Today, the conservatives are on the other side. Their preferred solution, articulated most comprehensively in an article by American Enterprise Institute fellow and former FCC chief economist Tom Hazlett (SEE RESOURCES), is private ownership of spectrum. The argument traces its origins to Ronald Coase's seminal 1959 analysis of FCC spectrum policy, which was one basis for Coase's eventual Nobel Prize in economics. The conservatives see the open spectrum notion as a new form of government regulation, because government must define the rules governing equipment to ensure cooperation over unlicensed bands.

Free the airwaves!

But the open-spectrum argument is more than that. It's not new wine in old bottles; it's a whole new vineyard. Many private monopolies on slices of spectrum are better than one, but they still don't represent a free market. The open spectrum view is that instead of regulating or assigning rights to use spectrum bands, government should set simple rules for equipment, and allow technology to do the rest. As with the current FCC "Part 15" rules, equipment licensing could be streamlined, with a significant amount of self-certification. (Competitors would be sure to point out where another company had broken the rules.)

Government must be involved one way or another, either to define private spectrum rights or to define equipment licensing rules, so that shouldn't be a factor. The question is whether that regulation creates a market based on equipment, or one based on infrastructure investment.

It's the difference between the telephone industry and the computer industry. The government gave AT&T a monopoly so that it had incentives to build out a nationwide network. AT&T developed the equipment and rented phones to customers. This arrangement produced many benefits, but rapid service introduction and price reductions weren't among them. The computer industry, by contrast, grew through cutthroat competition. The result was less uniformity, but more innovation.

By strengthening incentives to maintain current arrangements, privatization may prevent an efficient competitive market based on unlicensed spectrum. Notes Reed: "Just as land lots, once subdivided and sold to a lot of independent owners, are quite difficult to reassemble for a new kind of use, so inefficiently subdivided 'spectrum rights' will have a very high transaction cost whenever innovation might generate a huge benefit. The problem is that one 'defector' can block efforts to reorganize."

Commercial Meshed Wireless Networks

Monster mesh

As we've noted, 802.11 uses unlicensed spectrum but doesn't implement the fully decentralized, cooperative architecture that open-spectrum proponents advocate. 802.11 access points still must hook into traditional Internet backbones to deliver wide-area connectivity. A wireless network that fully realized the open-spectrum vision would be completely distributed. It would use a meshed rather than a star

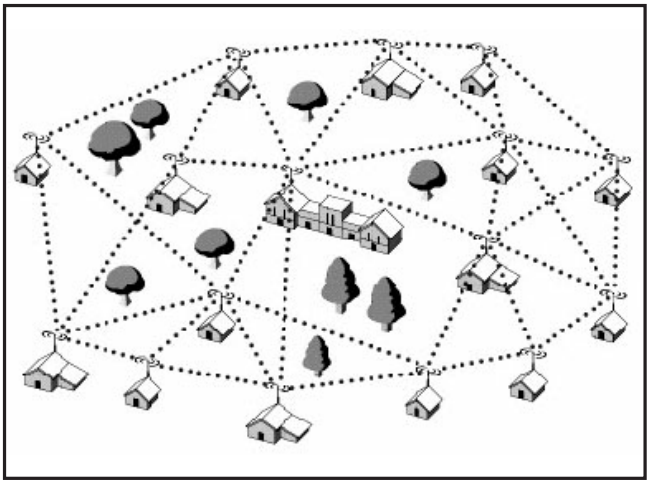

FIGURE 2: *Meshed wireless network (Illustration courtesy of Nokia.)*

topology for routing traffic. Every new node would extend and enhance the system, what's known as ad hoc networking.

In particular, mesh routing addresses a problem that has bedeviled most broadband fixed-wireless systems: line of sight. At the available frequencies, transmitters and receivers must be able to see one another. Hills, tree cover, intervening buildings and weather pose significant barriers in many residential areas.

In a meshed network, explains Lauren Hipp, head of sales and marketing for Nokia's wireless broadband routing line, "You only need line of sight to neighboring nodes in the network, not all the way back to the base station." Because the radios are close to ground level, high trees and other obstructions above them can function as a canopy that helps signal propagation.

Meshed wireless systems exist. They are commercially available today. And, among others, they are being sold by Nokia, the largest global vendor of wireless equipment. Nokia's Rooftop system and the competing Packethop and MeshNetworks technologies (SEE BELOW) all evolved from DARPA-funded research done by a group of organizations including the University of California at Santa Cruz and SRI International. In addition to commercial and military applications, the researchers envisioned meshed routing as a mechanism for shared-access community networks, akin to the grassroots 802.11 access points described above (SEE RESOURCES).

Nokia shouts from the rooftops

In September 1999, Nokia acquired Rooftop Communications, which had been formed in 1995 to commercialize wireless meshed routing technology. Rooftop marries standard 802.11 radios with a proprietary operating system, the Air OS, that handles channel scheduling and intelligent routing. Each node keeps a realtime map of network topology. When a subscriber joins or leaves the network, or spring foliage obscures what was previously a clear view between nodes, the devices automatically switch to the best possible routes. Nokia uses its own media access control (MAC) protocols for more robust performance, because the standard 802.11 ad hoc mode wasn't intended for a meshed network involving multiple hops.

Nokia's Hipp says the economic appeal of the technology for services providers is straightforward: "It makes it easy to deploy a network and get very high coverage right off the bat, without having to deploy a lot of infrastructure." Each wireless router simultaneously comprises a network access point, customer premises modem, firewall and home-networking gateway. The units list for about $1000, less in volume. In a typical residential network, the system supports connection speeds of 384 to 512 kbps, comparable to DSL and cable modems, with somewhat faster speeds possible with lower user densities.

Nokia is in discussions with various carriers and ISPs, and is also selling Nokia Rooftop systems via Tessco Technologies, a wireless equipment reseller. Its first two commercial deployments are with Advanced TelCom Group, an integrated communications provider based in Santa Rosa, CA, and Meer.net, an ISP in Mountain View, CA. Nokia sees the US as the primary market today, because demand is higher and because US rules allow 2.4-GHz systems to operate over greater distances than in Europe, making meshed networks more economical.

Other meshed systems

Nokia is furthest along in rolling out a meshed system, but it has competition.

SRI International, the storied Silicon Valley R&D organization, has been doing wireless TCP/IP networking since 1977. Its Packethop meshed routing software is radio-independent, so it can operate on 802.11, Bluetooth or licensed frequencies including 3G and multipoint microwave distribution service (MMDS). SRI claims its solution offers better efficiency and scalability than alternatives, and is more flexible for deployment in different application scenarios. The initial Packethop licensee is SpeedCom Wireless, a publicly traded wireless product vendor.

Meanwhile, MeshNetworks, based in Maitland, FL, has commercialized mesh routing technology used by the military. The company was founded in January 2000 with $27 million in funding from Patricof & Co., Redwood Ventures, ITT, 3Com, Apax Partners and BancBoston. It claims up to 2 Mbps sustained data rates per subscriber, with quality of service to support voice as well as data and video transmissions. MeshNetworks sells technology licensed, hardware and chips to OEMs and operators. It recently signed an agreement with Itochu, a large Japanese manufacturing company, to develop and market its technology in Asia.

UltraDevices: bottom-up wireless

The most interesting wireless meshed routing startup is UltraDevices. The company has built prototype 802.11a meshed-routing units which it plans to begin trials with potential customers before the end of the year, with the goal of commercial deployments in mid-2002.

UltraDevices chairman Robert Berger, who founded the ISP Internex and sold it to Concentric Networks in 1997, explains that 802.11 devices have commoditized radio frequency data communications, creating an opportunity that wasn't available to previous fixed wireless systems such as AT&T's recently cancelled Project Angel. "We believe that, using this as the baseline, we can hit a price point that hasn't been able to be hit elsewhere in the wireless outdoor market," says Berger.

ULTRADEVICES INFO

Headquarters: Mountain View, CA
Founded: August 2000
Employees: 6
Funding: $1.8 million from several Silicon Valley networking veterans
URL: www.ultradevices.com

UltraDevices hopes to sell its devices for $500 in volume, more than cable or DSL modems but significantly less in total cost when you factor in the additional infrastructure wireline systems require. That's substantially less than Nokia's system goes for today, though everyone's prices will drop over time. Berger believes in taking advantage of the full 802.11a standard and layering mesh routing software on top. What this approach lacks in initial capabilities it will make up for in price and faster innovation over time, Berger says, just as Ethernet did for wired LANs. "The Ethernet philosophy always wins," he argues.

That bottom-up philosophy carries over to UltraDevices' deployment vision. The company's initial customers will be entrepreneurial ISPs and "micro-carriers" looking for ways to route around the established cable and telco last-mile infrastructure. "The holes [in cable and DSL deployment] are really craters, especially once you get out of the first-tier markets," Berger says. "Lots of America has no broadband." Small providers know local markets, and they don't need many subscribers to make their business models work.

This guerrilla approach contrasts with the failed broad-scale rollout plans of competitive DSL providers and wireless ISPs such as Metricom. At some point, aggregators will tie together these local networks, as was the case for dial-up ISPs. UltraDevices' product can be deployed in a dual-radio configuration, using 802.11a for the meshing between nodes and an 802.11b transmitter to create local hotspots around each node in a heavily trafficked area.

Spectrum Futures

Open spectrum is a philosophy that opens up possibilities. As David Reed explains: "There are a lot of 'real options' out there for using radio in new ways and for new uses. Creating a market for such options may be much better than creating a market for spectrum – just as creating a market for corn futures is probably better than trusting in a market for farmland to decide how much to allocate to corn."

This view is consistent with a larger challenge to conventional networking models that we explored in our recent broadband report (SEE **RELEASE 1.0**, JULY/AUGUST 2001). Perhaps broadband connectivity is more like the automobile industry than the highway system. Perhaps it's more efficient for end-users to buy their broadband hookups the way they buy PCs, instead of renting network services the way they pay for electricity. [FOR MORE, SEE DAVID ISENBERG'S "CUSTOMER OWNED NETWORKS" ARTICLE IN RESOURCES.] As municipalities explore such models for "condominium fiber," open-spectrum proponents are proposing something similar for wireless services.

The proliferation of 802.11 access points hints at this scenario. The potential would be even more exciting, however, if there were more spectrum available for shared use, without 802.11's limitations and the congestion in the 2.4-GHz band.

Hardware to the rescue?

A credible scenario for overcoming the government bottleneck involves a group that traditionally has followed the consequences of spectrum allocation decisions rather than trying to change the system: hardware vendors.

Hardware vendors stand to benefit in an open-spectrum world, because they will no longer depend on a few network operators to dictate what equipment can be used. Says Yochai Benkler, "The most important political change will happen, if it happens, when the equipment manufacturers understand that with a few tweaks of the regulatory system, there's a whole new hardware market that opens up to them. They can become Detroit to the wireless services' railroad barons." The problem is that equipment vendors make their money today selling to the very wireless operators they would be competing with.

One possibility suggested by Dale Hatfield would be for a coalition of equipment manufacturers to jointly acquire some spectrum for unlicensed use, either directly through an FCC auction or in a secondary market transaction. They could create, in

effect, a private spectrum preserve. Equipment would have to conform to specifications that the vendor consortium defined, but there would be no central service providers controlling use of the spectrum. So far, though, no hardware company has expressed interest in this idea.

The near-term outlook for open spectrum is grim. The politics and the industry recognition simply aren't there today. But time is on the side of open spectrum. Wireless devices are only going to get smarter and cheaper. The danger now is that governments, by parceling out more licensed spectrum and establishing secondary-markets regimes that discourage shared usage, will inadvertently work against the direction of technology.

We have the 5-GHz U-NII band today because back in 1995, Apple brought a remarkably forward-thinking proposal to the FCC for wireless data networks in schools. A group of other technology vendors called WINForum joined the FCC process to advocate a variation of Apple's idea. Without that industry lobbying, there would be no 802.11a market, because 802.11a uses the U-NII spectrum. Who knows what future unlicensed services will take advantage of that band.

Intel last month filed comments with the FCC supporting ultra-wideband services and urging the Commission to "expeditiously adopt regulations" allowing the technology to go forward commercially. For the most part, though, regulators hear often from the companies with the most to lose from a spectrum commons, and seldom from those in the technology industry who thrive on innovation. Who will step up to the open-spectrum challenge? ■ R 1.0

New Media's Next Wave

BY KEVIN WERBACH

Convergence may happen sooner than you think. The stars are aligning for a new wave of interactive media services to gain traction in the US in 18-24 months. Most of them are being driven by established media companies rather than upstarts proclaiming a revolution, but the implications are just as significant.

Since the 1980s, it has been clear that video and audio entertainment would some day be delivered through digital networks, changing the distribution economics of the entire media industry (SEE **RELEASE 1.0**, NOVEMBER 1999). It has become equally clear that business plans based on that notion are prone to spectacular failure.

- In the 1990s, phone companies and cable operators poured billions into interactive television and video dialtone projects, only to see them flop.
- High-profile Internet-based media efforts such as DEN and Pseudo imploded with the collapse of the technology stock bubble.
- Microsoft, after shelving its high-profile interactive TV plans, paid $425 million for WebTV but has had trouble attracting subscribers.
- The record industry forced Napster to shut down its wildly popular music-sharing service (though it may rise again with the help of Bertelsmann).
- Personal video recorders (PVR) from Tivo and ReplayTV won investments from industry heavyweights and provoked anxious speculation about the death of TV, but in two years have garnered less than half a million users.

Here's why the story might end differently this time. End-user hardware is reaching the right maturity and price points. Video-on-demand services are finally going into production as major cable operators complete their digital upgrades. User interfaces are becoming easy enough for couch potatoes to use. Internet-based technologies to support economical video distribution to mass audiences are becoming available.

Hardware

PVR functionality, which lets viewers select, pause, fast-forward and store TV programming in digital format, is becoming a commodity. Tivo and Replay, the latter now part of SonicBlue, are releasing next-generation boxes. They are being joined in the interactive home entertainment race by Microsoft with its Ultimate TV, AOL

with AOLTV and Rearden Steel from WebTV creator Steve Perlman. CacheVision, a joint venture of Thomson Multimedia and Seagate, is licensing PVR software to cable operators, as is Replay. Residential gateway vendor Ucentric (SEE **RELEASE 1.0**, JULY/AUGUST 2001) has added PVR functionality to its offering. Hard drive prices continue to drop, making digital video storage roughly five times more affordable than it was when Tivo and Replay first launched.

While some vendors, including Intel and 3Com, have shut down their convergence appliance efforts, others such as Apple and HP are jumping in. Apple's iPod digital music player is generating great buzz and may be the start of a new device category.

New consumer-electronics devices always take several years to reach critical mass. The VCR and the CD player are two examples that are ubiquitous now but were curiosities for some time after they were introduced. PVRs, as standalone equipment or components of other devices, look to be approaching the upward slope of the growth hockey stick. Virtually everyone who has one finds it indispensable.

Software and programming

Video on demand (VOD) has always been high on the list of killer services for converged media. If users will pay to watch movies in theatres, on rented videotapes and on pay-per-view cable channels, surely they will pay for the ability to select any movie they want and see it immediately on their screen. For several reasons, proving this hypothesis has been difficult. But the industry hasn't stopped trying. Intertainer, which licensed rights to movies from major Hollywood studios, now offers on-demand video streaming in 35 markets in partnership with Microsoft's MSN. Most major cable operators are in the pilot phase for significant VOD rollouts.

The killer content for VOD may actually not be movies, which are widely available through other channels. Niche programming, such as how-to videos, health information, education and recreational video, is extremely cheap to acquire and doesn't make sense to deliver under conventional broadcast economics. For example, Chaos Media Networks in Santa Clara, CA, has licensed over 12,000 hours of specialty video programming which it packages into channels for cable operators. The company's software engine runs in set-top boxes and handles user interface, channel guides, advertising and e-commerce tie-ins.

COMING SOON

- Online collaboration tools.
- Distributed distribution.
- Patents pending.
- Real-time computing.

- And much more. . . (If you know of any good examples of the categories listed above, please let us know.)

User models

Another cause for optimism is the way companies are refining usage models. New forms of media always produce new user experiences. Videotapes changed the way people experienced movies, for example. For interactive services, "delayed media" interfaces that offer users free choice but don't fulfill those requests immediately seem to be catching on.

A good example is Kontiki, a peer-to-peer rich media distribution tool from a start-up led by ex-Netscape executive Mike Homer. With Kontiki, you select video or other files you want to download, and the system sends them to you in the background, with notification when the content is available on your hard drive. This may happen in a few seconds or overnight, depending on the content and your connection. There are similarities between Kontiki's interface and Tivo's, which becomes less surprising when you realize that Homer sits on Tivo's board.

The benefit for the services is that they don't have to dedicate the bandwidth and storage it would take to deliver any content at any time. Like the Internet's TCP/IP protocols, delayed media providers allow some unpredictability in order to make the system as a whole more robust.

Delayed media won't replace what we're used to, because some content has to be viewed immediately (such as breaking news footage) or at the same time others view it (such as the *West Wing* episode you'll discuss with co-workers tomorrow around the water cooler). Much of the time, though, the wait (or the savings from the wait) will be worth it.

Getting there

These developments don't all lead in the same direction. Some involve rich media delivered over the Internet to PCs, while others mean distribution of archival programming over private networks to TVs. Breaking through the advertising-driven one-way broadcast model of TV and radio will begin the experimentation with the next mass medium, not end it.

To use a familiar analogy, pen computing was a fundamentally solid idea despite all the failures it spawned. It's not clear that we've seen the Palm of converged media yet, but we're moving toward critical mass. Everything described above will be available by the first half of 2002. It will take some time to work out the bugs and grow user bases, suggesting that 2003 could be a very interesting year. ■ R 1.0

Resources & Contact Information

C.J. Cornell, **Chaos Media Networks**, 1 (408) 330-0955; cj@chaosmediaventures.com
Dewayne Hendricks, **Dandin Group**, 1 (510) 573-0561; fax, 1 (510) 573-0562; dewayne@dandin.com
Hong Chen, **GRIC**, 1 (408) 955-1111; hchen@gric.com
Clark Dong, **HereUare**, 1 (408) 551-0909; fax, 1 (408) 241-4698; clarkd@hereuare.com
Mike Homer, **Kontiki**, 1 (650) 625-3073; mike@kontiki.com
Tim Shepard, **MIT**, 1 (617) 253-6002; fax, 1 (617) 253-2673; shep@lcs.mit.edu
Rob Flickenger, Schuyler Erle, **NoCat.org**, 1 (707) 827-7151; rob@oreillynet.com; sderle@oreillynet.com
Lauren Hipp, Dave Beyer, **Nokia**, 1 (650) 625-2000; fax, 1 (650) 625-2380 lauren.hipp@nokia.com
Yochai Benkler, **New York University**, 1 (212) 998-6738; yochai.benkler@nyu.edu
Anthony Townsend, **NYCWireless**, 1 (212) 998-7500; fax, 1 (212) 995-3890; anthony.townsend@nyu.edu
Scott Seaton, **SRI International**, 1 (650) 859-6222; scott.seaton@sri.com
Lawrence Lessig, **Stanford Law School,** 1 (650) 736-0999; fax, 1 (650) 723-8440; lessig@pobox.com
Robert Berger, **UltraDevices**, 1 (408) 882-4755; fax, 1 (408) 490-2868; rberger@ultradevices.com
Dale Hatfield, **University of Colorado**, 1 (303) 492-6648; dale.hatfield@colorado.edu
Dan Lowden, **Wayport**, 1 (512) 519-6149; fax, 1 (512) 519-6150; dlowden@wayport.net
Scott Rafer, **WiFinder**, 1 (925) 385-0104; fax, 1 (925) 299-9298; rafer@wifinder.com
David Reed, 1 (508) 785-2675; fax, 1 (508) 785-8194; dpreed@reed.com

For further reading:

Tom Hazlett, "The Wireless Craze, The Unlimited Bandwidth Myth, The Spectrum Auction Faux Pas, and the Punchline to Ronald Coase's 'Big Joke'" AEI-Brookings Joint Center for Regulatory Studies working paper (January 2001) – http://www.aei.brookings.org/publications/working/working_01_02.pdf

Yochai Benkler, "Overcoming Agoraphobia: Building the Commons of the Digitally Networked Environment," 11 **Harvard Journal of Law and Technology** 287 (1998) – http://www.law.nyu.edu/benklery/agoraphobia.pdf

Tim Shepard, "Decentralized Channel Management in Scalable Multihop Spread-Spectrum Packet Radio Networks," MIT dissertation (1995) – ftp://ftp.lcs.mit.edu/pub/lcs-pubs/tr.outbox/MIT-LCS-TR-670.ps.gz

Communitywireless access groups – http://www.personaltelco.net/index.cgi/WirelessCommunities

Dave Beyer, Mark Vestrich and JJ Garcia-Luna-Aceves, "The Rooftop Community Network: Free, High-speed Access for Communities" (1999) – http://www.ksg.harvard.edu/iip/doeconf/beyer.html

David Isenberg, "SMART Letter #56: Era of Customer-Owned Networks" – (June 7, 2001) – http://www.isen.com/archives/010607.html

Calendar of High-Tech Events

2001

DECEMBER 3-5 **PRIVACY BY DESIGN** - Montreal, Canada. Zero Knowledge Systems brings together business, technology, and policy leaders. Call 1 (514) 286-2636 x7706; fax, 1 (514) 350-7587; email conference@zeroknowledge.com. More information at privacy.zeroknowledge.com/privacybydesign2001/registration.asp

DECEMBER 5-6 **AD:TECH** - New York, NY. Marketers gather to ponder the online advertising depression, and to look for lights at the end of the tunnel. Email Jenabeth Ferguson, jferguson@imark-com.com. www.ad-tech.com

DECEMBER 5-7 **GLOBAL CITIZEN NETWORKS 2001** - Buenos Aires, Argentina. The second World Congress of Citizen Networks gathers around the theme: "Renewing Communities in the Digital Era." For additional information, contact the secretary of Global CN2001 at secretariado@globalcn2001.org. www.globalcn2001.org

DECEMBER 10-13 **STREAMING MEDIA EAST 2001** - New York, NY. Touted as the world's largest Internet audio and video event. For more information, go to www.streamingmedia.com/east or call 1 (888) 301-8890 to register. www.streamingmedia.com/east

DECEMBER 10-13 **FALL INTERNET WORLD** - New York, NY. The mother of all Internet trade shows makes its east coast stop, after a two-month postponement due to the events of September 11. To register, call 1 (800) 632-5537; fax, 1 (203) 559-2814. www.internetworld.com

2002

JANUARY 7-11 **MACWORLD CONFERENCE & EXPO** - San Francisco, CA. The most comprehensive gathering for the Macintosh community. Register online, or call 1 (800) 645-EXPO for information. www.macworldexpo.com

JANUARY 8-11 **CONSUMER ELECTRONICS SHOW** - Las Vegas, NV. The place to be to see the latest, hottest consumer technology. Scheduled keynoters include Bill Gates, Bob Pittman and Carley Fiorena. Register online,email CESinfo@ce.org for help or call 1 (703) 907-7605. www.cesweb.org

MARCH 24-27 **PC FORUM** - Scottsdale, AZ. EDventure Holdings' premier conference, now in its 25th year. Look for a speaker list and more details soon on our Website at www.edventure.com/pcforum. Contact Daphne Kis with any questions at 1 (212) 924-8800 or daphne@edventure.com. E K

E Events Esther plans to attend.
K Events Kevin plans to attend.

Lack of a symbol is no indication of lack of merit. The full, current calendar is available on our Website, www.edventure.com. Please contact Irene Lawrence (irene@edventure.com) to let us know about other events we should include.

The 25th Annual PC Forum:
Back to the Frontier

March 24 to 27, 2002

The Fairmont Scottsdale Princess, Scottsdale, Arizona

For a quarter of a century, PC (Platforms for Communication) Forum has started fresh each year. We stay ahead of the changing times, picking out topics that shed light on the years to come. We bring together the leaders and the visionaries, the industry giants and the path-breaking startups... some of which grow into giants themselves! This coming March, for our silver anniversary, we will help the industry return to its pioneering spirit.

Join us at PC Forum to reinvigorate yourself for the challenges and opportunities ahead.

Please visit **www.edventure.com/pcforum.**

Release 1.0 Subscription Form

Join the other industry executives who rely on *Release 1.0* . You can also subscribe online at **www.release1-0.com.**

Your annual *Release 1.0* subscription costs **$795** per year (outside the US, Canada and Mexico, add **$55** for shipping), and includes both the print and electronic versions of 11 monthly issues; 25% off the cover price when you order from our online archives; a *Release 1.0* binder; the bound transcript of this year's PC Forum (a $300 value) and an invitation to next year's PC Forum.

NAME _____

TITLE _____ COMPANY _____

ADDRESS _____

CITY _____ STATE _____ ZIP _____ COUNTRY _____

TELEPHONE _____ FAX _____

EMAIL* _____ URL _____

**personal email address required for electronic access.*

☐ My colleagues should read Release 1.0, too!

Send me information about multiple copy subscriptions and electronic site licenses.

☐ Check enclosed ☐ Charge my (circle one): AMERICAN EXPRESS MASTER CARD VISA

CARD NUMBER _____ EXPIRATION DATE _____

NAME AND BILLING ADDRESS _____

SIGNATURE _____

Please fax this form to Natasha Felshman at 1 (212) 924-0240. Payment must be included with this form.

Your satisfaction is guaranteed or your money back.

If you wish to pay by check, please mail this form with payment to: EDventure Holdings, 104 Fifth Avenue, 20th Floor, New York, NY 10011, USA. If you have any questions, please call us at 1 (212) 924-8800; email us@edventure.com; www.edventure.com.